

Plan to Address the Mandated Ramp Down of Experiment Offline Production Operations Service (OPOS)

Tanya Levshina, Anna Mazzacane, Margaret Votava

[Background](#)

[Current Status](#)

[Revised goals based on need to ramp down OPOS](#)

[Proposal for Initial Reassignment of Activities and Effort](#)

[Organization Changes in Scientific Computing Services \(SCS\) Quadrant in Support of these Goals](#)

[Ramp Down Plan](#)

[Appendix I](#)

[Appendix II](#)

Version History				
Version	Date	Author(s)	Approved by Change Number	Change Summary
1.1	03-30-2016	Tanya Levshina/Anna Mazzacane		Draft initial document

Background

This proposal is in response to the following changes:

- The large and increasing experiment requests for software and computing resources, identified during SCPMT
- The budget constraints faced by the Scientific Computing Division(SCD).

As a result of these circumstances, the head of SCD, Panagiotis Spentzouris, has asked SDSC to propose a path to ramp down of the Experiment Production Operations service area/offering

and to plan for re-assignment of the effort to existing high priority gaps in resources within the department. (See the scientific portfolio management team (SC PMT) report and SCD response [here](#)).

Current Status

During the past year OPOS has helped the experiments (Minos, MINERvA, NOvA, DUNE) in organizing and streamlining their production workflow, improving version control, monitoring and debugging failure of jobs and data handling on Fermi resources, on the OSG and HEP Cloud. OPOS worked on providing more secure environment for production jobs submission by working with security group on the policy for using shared accounts and distributing service certificates. OPOS also contributed to requirements gathering for POMS.

OPOS efforts to make off-line production for experiments well organized, with supporting processes and documented procedures have increased the productivity and robustness of the experiment production jobs for those experiments supported.

Revised goals based on need to ramp down OPOS

Our goal is to have experiments continue to benefit from the past years of work of OPOS and to broaden benefit to all experiments supported by FIFE. We thus are adding to two thrusts made possible by dropping the actual jobs submission activities:

- Incorporate the good practices that have been developed by OPOS into the “on-boarding to OSG” and user support processes for all experiments supported by FIFE . This leverages work already started in OPOS to fill in the gaps in and improve user support for experiments using FIFE and the OSG.
- Create a framework for the training, consulting and mentoring of experiment post-docs who are taking over the production job submissions, providing a bidirectional channel for the promulgation of the good practices to the experiments and continuous communication from experiments to FIFE user support of production activities and needs as they occur .

Proposal for Initial Reassignment of Activities and Effort

1. Extend LArSoft specific Continuous Integration (CI) Service: OPOS (and FIFE support) has encountered several/many instances where experiment code being used for production running has bugs and problems that result in wasted computing and human resources. We strongly recommend that experiments improve their practices in testing their workflow thoroughly before submitting to production running. Automated build and test of software provides feedback to experiments as soon as possible in a controlled environment and helps identify problems introduced by code changes, thus avoiding

waste of resources. As tests are built up more such issues can be found and addressed. The LArSoft CI framework was designed and implemented to be general and will, with some configuration and modifications, support other experiments in test their workflows. (0.3 FTE + 1 intern/associate)

2. Contribute - effort, documentation, resulting good practices - to application profiling: This would start with the current contributions to the memory profiling activity for MicroBooNE LArSoft applications and then ensure that the processes developed can be re-used both for LArSoft applications and for other experiments. (0.3 FTE)
 3. On-board all experiments to use the Production Operations Management Service (POMS) developed by SCS for OPOS: Several experiments have already expressed interest to use POMS for production submission and troubleshooting. In addition we propose to increase our contributions to satisfying POMS feature and improvement requests as well as troubleshooting its use and providing better user documentation. (0.3 FTE + 1 intern/associate)
 4. Build on the existing OPOS support tools and services to help support more efficient, robust job and data workflows for FIFE experiments. (0.2 FTE + 1 intern/associate)
 5. Anna Mazzacane will approach the experiments to discuss how to provide a framework to train and mentor experiment personnel who will be submitting and tracking production jobs once OPOS ramps down. She will develop a plan on how this will work for both the experiments and the FIFE user support group and present that to Division and Experiment management for input and support. The goal will be to pilot the idea with one or two of the legacy experiments in the summer and fall of 2016. (0.5 FTE)
 6. OPOS has done significant amount of work improving NOvA production submission and absolutely crucial for AWS submission and monitoring. (0.05 FTE)
- (total is 1.7 FTE + 3 interns/associates)

Organization Changes in Scientific Computing Services (SCS) Quadrant in Support of these Goals

1. Projects
 - Experiment Continuous Integration framework :Move project management to the User Support for Distributed Computing (USDC) group and add effort; with major development remaining in the Scientific Data Processing Solutions (SDPS) Department.
 - POMS: Move project management and support for experiments to USDC; with major development and operational support remaining in SDPS.

2. Operation and support:

- Service Certificates for shared accounts handling support
- Extended support for “best practice” for production job submission to be in the USDC group instead of OPOS.

3. Groups and Personnel

- Transfer OPOS members to USDC group led by Joe Boyd:
 - Anna Mazzacane - Make Deputy head of USDC, CI - project management, POMS - project management and deployment.
 - Vito Di Benedetto - CI technical contributor; Additional contributions to application profiling.
 - Felipe Alba (Intern 1: leaving summer 2017) - experiments on-boarding (GPgrid, OSG, HepCloud) , MINOS & Minerva ramp down
 - Paola Buitrago (Intern 2: leaving on July 1st 2016) - POMS development, NOVA HEP Cloud, NOvA ramp down
 - Taylor Propp (on call) - experiments production jobs monitoring
 - Michele Fattoruso (Intern 3: arriving April 4th 2016, leaving spring 2017?) - experiments on-boarding (GPgrid, OSG, HepCloud) , POMS development (?)
- Transfer OSG visiting scientist to USDC group (only needed if re-organization happens before July)
 - Juan Morales (leaving on July 1st) - GratiaWeb for Gratia and GRACC
- Replace Interns when they leave with associates (A2-3)
 - 2016: Replace Intern 2 with A2 Associate: CI development and CI experiment on-boarding , documentation
 - 2017 : Replace intern 1 with A2 Associate: Monitoring and debugging
- CMS Production team will also be moved to USDC.
 - Jen Adelman-McCarthy
 - Gaston, Eliana, Jorge - Interns

Ramp Down Plan

1. Meet with experiments that are currently are using OPOS services and brief them on the upcoming changes.
2. Identify post-docs who will be responsible for running production campaigns for NOvA, MINOS, MINERvA, DUNE and LArIAT.
3. Start training these post-docs to submit and troubleshoot jobs via POMS and use FIFEMON for monitoring.
4. Gradually decrease OPOS involvement in job submission. Set up a deadline to discontinue OPOS support for each experiment.

5. Anna becomes POMS project manager - late Spring 2016
6. Transfer Anna and Vito to USDC. Anna - Deputy group lead (Summer 2016)
7. Decrease Vito's involvement in production submission and end it by summer 2016
8. Anna becomes CI project manager - Summer 2016
9. Vito should be working on CI for LArSoft and other experiment by September 2016. He will start working CI monitoring as well.
10. Transfer two remaining interns (Felipe and Michele) to USDC group and assign them to CI and monitoring program. (Summer 2016?)
11. Train Lisa to handle all issues related to CMS visiting scientists (interviewing, visa, housing, etc).
12. Transfer CMS group under USDC
13. Hire associate for CI and Monitoring(!!!! - we need it) - late summer 2016
14. Continue to do AWS submission until we have a robust procedure with all the checks and alarms (will be done by one of the visiting scientists)

Appendix I

DRAFT for associate position for monitoring.

Computing/IT Associate

The Scientific Distributed Computing Solutions department is looking for a talented, high-motivated and independent individual to contribute to the work program of the Scientific Computing Division. The department enables Fermilab experiments, projects and partners to carry out their science on distributed computing resources, including the Fermilab facilities, the Open Science Grid, private and commercial Clouds. This is an entry-level position and will require learning, supporting, and expanding the next generation monitoring and accounting service at Fermilab. We are monitoring and tracking jobs running on thousands of cpus and moving Terabytes of data at Fermilab and across the US. Customers of the tools include scientists, facility operators, and funding agencies.

Responsibilities:

- Installs, configures, maintains and upgrade Elasticsearch instances in support of scientific processing requirements.
- Monitors instances and notify others of any issues.
- Work with developers to test and evaluate new grid accounting and monitoring dashboards.
- Assists in creating documentation for monitoring and accounting services.
- Participates in design and development of new dashboards, probes and reports.
- Communicates and coordinates with other people working closely on the project.
- Interacts with end users and supporting services as needed.

Qualifications and Essential Job Functions

- BA,BS in Computing or equivalent experience
- 0 - 2 years of experience in IT related fields
- Excellent verbal and communication skills
- A willingness to be challenged and learn new technology

Appendix II

Current assignments

Name	VO	Assignment	End of Contract	FY17
Filipe	NOvA FD	Stage data from tape to dCache (monthly)	July, 21st 2017	Transfer to the experiment
		MC production (on demand)		Transfer to the experiment
		Backprocessing (on demand)		Transfer to the experiment
	MINERvA	MC production (on demand)		Transfer to the experiment
	MINOS+	Keepup for MINOS+ (daily)		Transfer to the experiment
Paola	NOvA	raw2root keep-up (daily) about 2K jobs, 50K files	July 31st, 2016	Transfer to the experiment
		ND MC Reconstruction on AWS - on demand		Transfer to the experiment
	All	POMS feedback		Transfer to USDC
	All	AWS submission documentation		Transfer to the experiment
Qiulan	NOvA	reconstruction keep-up (daily) 400 jobs	April 1st, 2016	Felipe first, then transfer to the experiment
	All	Contribution to Monitoring: condor event log and dCache log parsing		Transfer to USDC

Tyler	NOvA	Production Systems Monitoring (daily)	June 30th, 2016	Transfer to the experiment
	MINERvA	OFF-line systems monitoring (daily)		Transfer to the experiment
Vito	NOvA	Monitoring the memory and disk usage of the raw2root and reco keepup processings (daily)	Permanent position	Transfer to the experiment
		Daily monitoring and scan of the offsites accepting jobs.		Transfer to USDC
	All	Developing the unified submission tool set		Transfer to USDC
		OPOS weekly meeting reports. Weekly SCS ops dashboard update		Transfer to USDC (ops report)
Michele			Arrives on April 4th, 2016	Will pick up Paola's assignments.
	DUNE			
	LArIAT	raw2digit		